

CANONICAL



Ceph Operations

Training Curriculum

Prepared For:

Canonical

Prepared By:

Radu Popescu, Mihai Gheorghe

Prepared On:

1 April 2020

Version:

1.0

Introduction & Scope

This course is intended to enable its attendees to assume full command of their on-premises Ceph cluster. The target audience is system administrators, DevOps engineers, and infrastructure operators. The main focus of the curriculum is understanding Ceph and all of its components, installing and configuring them, troubleshooting for each of the components, optimising the performance of the Ceph cluster, and troubleshooting different issues. Each chapter contains a theory session (presented by the instructor) followed immediately by a practical lab. The practical lab follows the main points described during the slide presentation and will provide full hands-on experience in working with Ceph. By combining the theoretical background with hands-on exercises, the course encourages better assimilation of the material and deeper understanding. Finally, the students will be provided with complete step-by-step documentation, ensuring that the material is readily available for future review.

The duration of the course is three days.

Agenda

1. What is Ceph

- Description of Ceph as software defined storage (SDS)

2. What is RADOS

- Description of RADOS architecture and functionality

3.1 Ceph frontends to RADOS cluster

- In-depth look at each Ceph frontend service:
 - RBD (block storage)
 - RadosGW (object storage)
 - CephFS (file storage)

3.2 Ceph daemons

- Description of each daemon running inside a Ceph cluster
 - Monitors
 - Managers
 - OSDs

- MDSs

3.3 Deploying Ceph cluster

- Prerequisites required for a Ceph cluster deployment
Lab 3A - installing required prerequisites
- Description of ceph-deploy tool
- ceph-deploy usage to install all Ceph components
Lab 3B - deploy Ceph components on all nodes

4.1 High-level operations

- Operating a cluster
 - service management
- Health of a cluster
 - OSD, PG, Crusmap states
- Monitoring a cluster
 - cluster status
 - cluster usage
 - monitor health-checks
 - checking service status
 - monitoring OSDs and PG

4.2 Low-level operations

- adding and removing OSDs - both manual and ceph-deploy tool instructions
- adding and removing MONs - both manual and ceph-deploy tool instructions

4.3 Data placement

- In-depth look at pools, placement groups and crush maps
- Pools administration:
 - description
 - specific commands
 - setting quotas
 - show usage
 - setting/retrieving custom parameter values**Lab 4A** - creating pools, setting replication size, quotas
 - erasure coded pools - description, profiles**Lab 4B** - creating erasure coded pool profiles, creating pools
- Placement groups:
 - description and architecture

- scrubbing and recovery
- Lab 4C** - setting PG number in pools, show PG maps for specific PG, query PG for information
- Crush maps:
 - definition
 - structure
 - devices
 - map types and buckets
 - rules
 - weight sets
 - modifying the crush map
 - crush tunables
- Lab 4D** - modifying crush map using "ceph osd crush"
- Lab 4E** - manually modifying crush map

5. Ceph RBD

- What is Ceph RBD
- Basic commands and usage
- Operations - snapshots, mirroring, live-migration
- Lab 5A** - creating RBD pools, working with snapshots, clones
- Lab 5B** - mounting RBD volumes to a server, modifying features available
- CephX authentication and authorization
- User management
- Lab 5C** - managing users, permissions; testing

6. CephFS

- What is CephFS
- Adding a MDS server
- Removing a MDS server
- Creating a CephFS filesystem
- Lab 6A** - creating a CephFS filesystem
- Mounting a CephFS filesystem
- Lab 6B** - mounting a CephFS filesystem
- Quotas
- Client capabilities
- Exporting CephFS
- User management
- Lab 6C** - working with users, client capabilities and quotas

7. Rados Gateway

- What is Rados Gateway
- Requirements
- Rados S3 API
- Rados Swift API
- User management
- Key management
- Capabilities management
 - Lab 7A** - working with users, keys and capabilities
- Quota management - user and bucket scope, global settings
- Usage - show and trim
 - Lab 7B** - quota management and usage statistics
- Bucket sharding
 - Lab 7C** - working with bucket sharding

8. Ceph Dashboard

- Description of Ceph Dashboard functionality
 - Lab 8** - installing and accessing Ceph Dashboard

9. Ceph performance tuning

- Hardware layer
- O/S layer
- Bluestore
- Failure domains
- Benchmarking Ceph

10. Troubleshooting

- Logs
- Cluster status
- Data scrubbing
- MONs troubleshooting
- OSDs troubleshooting
- PGs troubleshooting
- Rados Gateway
 - Lab 10A** - checking Ceph logs, modifying running configuration
 - Lab 10B** - fault scenarios